

Coordinate-Space Hartree-Fock-Bogoliubov Solvers for Superfluid Fermi Systems in Large Boxes

J C Pei^{1,2}, G I Fann³, R J Harrison^{3,4}, W Nazarewicz^{1,2,5}, J Hill³, D Galindo³, J Jia³

¹Department of Physics and Astronomy, University of Tennessee, Knoxville, TN 37996, USA

²Physics Division, Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, TN 37831, USA

³Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37830, USA

⁴Department of Chemistry, University of Tennessee, Knoxville, TN 37996 USA

⁵Institute of Theoretical Physics, Warsaw University, ul. Hoża 69, 00-681 Warsaw, Poland

E-mail: peij@ornl.gov

Abstract. The self-consistent Hartree-Fock-Bogoliubov problem in large boxes can be solved accurately in the coordinate space with the recently developed solvers HFB-AX (2D) and MADNESS-HFB (3D). This is essential for the description of superfluid Fermi systems with complicated topologies and significant spatial extend, such as fissioning nuclei, weakly-bound nuclei, nuclear matter in the neutron star crust, and ultracold Fermi atoms in elongated traps. The HFB-AX solver based on B-spline techniques uses a hybrid MPI and OpenMP programming model for parallel computation for distributed parallel computation, within a node multi-threaded LAPACK and BLAS libraries are used to further enable parallel calculations of large eigensystems. The MADNESS-HFB solver uses a novel multi-resolution analysis based adaptive pseudo-spectral techniques to enable fully parallel 3D calculations of very large systems. In this work we present benchmark results for HFB-AX and MADNESS-HFB on ultracold trapped fermions.

1. Introduction

The Hartree-Fock-Bogoliubov (HFB) equation of the Density Functional Theory (DFT) is suitable for describing superfluid Fermi systems by properly accounting for the self-consistent coupling between the particle-hole and particle-particle mean fields. Recently, we solved HFB equations in complicated geometries to describe nuclei and ultra-cold polarized Fermi gases [1]. The general HFB equation for a polarized system can be written as:

$$\begin{pmatrix} h_a(\mathbf{r}) - \lambda_a & \Delta(\mathbf{r}) \\ \Delta^*(\mathbf{r}) & -h_b(\mathbf{r}) + \lambda_b \end{pmatrix} \begin{pmatrix} u(\mathbf{r}) \\ v(\mathbf{r}) \end{pmatrix} = E \begin{pmatrix} u(\mathbf{r}) \\ v(\mathbf{r}) \end{pmatrix}, \quad (1)$$

where h_a and h_b are Hartree-Fock Hamiltonians for the two spin components, λ_a and λ_b are the corresponding chemical potentials, and Δ is the pairing potential.

In some cases, solving the HFB equation in large boxes in coordinate space is essential. Weakly-bound nuclei and the large-amplitude nuclear collective motion such as fission and fusion are examples of problems that require large-box calculations [2]. The ultra-cold atoms trapped in elongated optical traps also require a description involving large spatial dimensions [1], as well

as HFB description of nucleonic pasta phases in the neutron star crusts [3]. All these problems are both interesting and important but the underlying calculations are challenging.

Many HFB solvers used in the nuclear physics context are based on the basis expansion method employing harmonic oscillator wave functions. The configuration-space method is efficient but offers a fairly poor accuracy for cases involving weakly-bound systems and large deformations, see discussion in Ref. [2]. On the other hand, solving HFB equations directly in coordinate-space can offer very precise results. Unfortunately, because of numerical challenges involved, HFB calculations in non-spherical geometries require high computation cost. In this context, multi-core processor architectures such as the Jaguar and Kraken Cray XT5 with 12 cores per node, or Hopper Cray XE6 supercomputer with 24 cores per node, promise to revolutionize the deformed HFB problem. To fully take advantage of unique computational capabilities, HFB solvers should be adapted and improved in terms of scalability.

To this end, we developed two parallel HFB codes: 2D HFB-AX for axially symmetric systems and MADNESS-HFB for fully 3D systems. Both solvers take advantage of modern multi-core architectures. The following sections, overview computational and numerical techniques employed by these two codes and present benchmark calculations for ultra-cold Fermi atoms. To illustrate the deformed HFB problem in coordinate space, we present an extreme application to novel pairing phases in polarized cold Fermi gases in extremely elongated traps [4].

2. Two-dimensional HFB solver HFB-AX

In HFB-AX [2], the wave functions are discretized on a 2-D grid (r_α, z_β) with the M -order B-splines:

$$\psi_{n\Omega\pi}(r_\alpha, z_\beta) = \sum_{i,j} B_i^M(r_\alpha) B_j^M(z_\beta) C_{n\Omega\pi}^{ij}, \quad (2)$$

where the coefficients C_n^{ij} can be obtained by diagonalizing the Hamiltonian matrix. To reduce

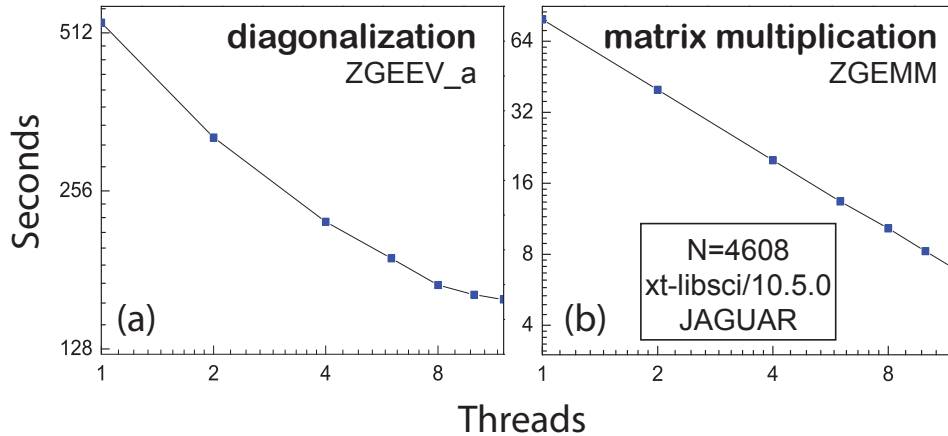


Figure 1. Left: performance of the complex eigensolver ZGEEV_a using different number of threads (ZGEEV_a is a modified version of LAPACK's ZGEEV). Right: performance of the complex matrix multiplication routine ZGEMM. The two routines are employed by HFB-AX and tested on Jaguar with the xt-libsci/10.5.0 library.

computational costs, we assume reflection symmetry; hence, the intrinsic parity π is a good quantum number that can be used to enumerate HFB eigenstates. Another quantum number preserved by HFB-AX is the angular momentum projection Ω on the symmetric axis. The boundary conditions are implemented through first- and second-order derivative operators. The

precision of computations depends on the mesh size, order of B-splines, and the box size. Those parameters determine the numerical effort involved.

In HFB-AX, the diagonalization of the complex Hamiltonian matrix takes the bulk of the computing time. Diagonalization blocks of given Ω and π quantum numbers are assigned to computing cores through MPI communications. We modified the LAPACK diagonalization routine ZGEEV to calculate only the selected eigenvectors. For calculations of a heavy nucleus with the mesh size of 0.6 fm, 10-th order B-splines, and a box size of 24 fm, the rank of Hamilton matrix is $N=4608$. To take the advantage of multi-core architectures, we use the threaded LAPACK library, which can speed up the diagonalization by a factor of 3 with 6 threads, as shown in Fig. 1(a). For the matrix multiplication problem in Fig. 1(b), the routine ZGEMM scales perfectly. For other matrix operations, we use OpenMP for parallel calculations. The modified Broyden method is used to accelerate the iteration convergence [2]. The hybrid MPI and OpenMP programming have made HFB-AX an accurate and fast HFB solver for large box calculations. Besides the threaded library on Cray machines, one can also utilize the threaded AMD's ACML library or the threaded Intel MKL library. In the near future, the multi-core+GPU Cray XT6 machine is expected, and the corresponding MAGMA library [5] can further improve the capability of HFB-AX.

The precision of HFB-AX has been demonstrated in calculations of weakly-bound nuclei in which the coupling to scattering continuum is essential [2, 6]. In large-box calculations, very dense quasi-particle spectrum is obtained, and the non-resonant continuum contribution can be precisely taken into account by means of the direct integration [6]. Furthermore, HFB-AX is expected to provide precise solutions for the continuum QRPA description of excited states by means of the recently developed finite-amplitude method [7]. The recent success of HFB-AX is the prediction of rapidly oscillating pairing potential in highly elongated traps [4] by using the SLDA [8] and ASLDA [9] energy density functionals for polarized Fermi gases.

3. 3D MADNESS-HFB Solver

The software MADNESS-HFB [10, 11] is a fast $O(N \log \epsilon)$ method based on the multiresolution analysis and low separation rank methods of representing and approximating functions and operators for solving Schrödinger and Lippman-Schwinger equations up to user-determined precision ϵ . MADNESS is an acronym for Multiresolution Adaptive Numerical Environment for Scientific Simulations [12]. The application of multiresolution analysis separates the behavior of functions and operators at different length scales in a functional way. Interesting mathematical and computational feature in our MADNESS library implementation is that each of the operators and wave functions has its own adaptive structure of refinement to achieve and guarantee the desired accuracy.

The representation of functions and operators in MADNESS is based on an adaptive pseudo-spectral representations of functions and operators using the discontinuous Alpert's multiwavelets [13, 14] and the low-separation rank representations and applications of Green's functions. Alpert's multiwavelets have compact support. For representing functions in 3D, we use tensor products of 1D multiwavelets. The two-scale relations provides for an adaptive approximation of the expansion in terms of multiwavelet basis as a function of accuracy in terms of the regularity of the solutions. The support of the bases functions in each level resembles that of structure adaptive refinement. A schematic figure of such a refinement process is shown in Fig. 2. Using two-scale relations, the multiwavelet basis provides a way of constructing an adaptive pseudo-spectral method. The multiwavelet basis is a local basis of degree k defined on the interval $(0, 1)$ generated from scaling functions which are defined by shifted, rescaled and orthonormalized Legendre polynomials from degree 0 to $k - 1$ from $(-1, 1)$ with a value of 0 defined outside of $(0, 1)$. The subspace, of $L_2((0, 1))$, spanned by scaling functions of degree k is denoted by $V^k = \phi_k(x) = \sqrt{k+1}/2 P_i(2x-1)$ on $(0, 1)$ where $P_i(x)$ is the i -th Legendre




Figure 2. A schematic picture illustrating the two-scale generated refinement process of a multi-resolution representation of a function in 2D with 3-levels.

polynomial on $(-1, 1)$. For each level n , we further define an ascending sequence of the subspaces $V_n^k = \phi_{jl}^n(x) = 2^{n/2} \phi_j(2^n x - l)$, $j = 0, \dots, k-1$, $l = 0, \dots, 2^n - 1$. Specifically, $V_k^0 \subset V_k^1$; the multiwavelets basis, denoted by W_k^0 , is the orthogonal complement of V_k^0 in V_k^1 . For each level n , we can further construct W_k^n in the same way of shifting and rescaling the basis functions of W_k^0 as in V_k^n to construct W_k^n . Thus,

$$\begin{aligned} V_0^k &\subset V_1^k \subset V_2^k \subset \dots \subset V_n^k, \\ V_n^k &= V_0^k + W_0^k + W_1^k + \dots + W_{n-1}^k. \end{aligned} \quad (3)$$

The notable features of the multiwavelet bases are that at each level the support of the basis functions are of width 2^{-n} , and that there is an exact algebraic relationship between the basis functions at level n and level $n+1$. In particular, the coefficients of a representation of a smooth function in the multiwavelet basis decays proportionally to the width of the interval at level n . Thus, we can estimate the accuracy of the representation and truncate coefficients below a specified threshold for a sparse representation.

In 3D, we use a tensor product basis generated from the 1D multiwavelet basis; the representation of operators uses a non-standard approximation or a low separation rank approximation based on Gaussians. The structure of the support of a function resembles that of an oct-tree. Each node of the oct-tree consists of a tensor of coefficients. The truncation of the sparse coefficients produce a pruned tree. The oct-tree is distributed across the nodes of a massively parallel computer referenced by a global hash-table for each node of the tree. The singular integral operators are represented using a low-separation rank approximation of the Green's functions [15] with details described in [16].

The MADNESS library uses a combined MPI and pthread parallel computing method in a task based computing model with a task graph scheduling and queue on each node. On each node of a parallel multicore computer, one core is devoted to processing internode communication

via MPI, and one core is devoted to handling thread scheduling and task allocation and queue scheduling within the node. The code is polymorphic with the use of C++ templates. Each operation on functions and the application of operators are defined in terms of tasks to be performed on the nodes or tensors of the “oct-trees” in the representation the functions or operators. In addition, a data path and flow path dependency analysis is performed to determine when different steps of the code can be overlapped and scheduled in each node’s task queue so that distributed multi-threading computation can be performed. Furthermore, since different functions and operations have different data and work-loads, a user directed dynamic load-balancing of data memory can be performed to permit more efficient data and work load distribution.

The solution methodology for the MADNESS-HFB consists of two main steps: (i) diagonalization of the Hamiltonian matrix using wave functions expanded in the multiwavelet bases in 3D, and (ii) computation of each of the wave functions using the associated Lippman-Schwinger equation. Recall that the Green’s function for the operator $-\Delta - \lambda$ is $e^{-\lambda r}/r$, the Yukawa potential. The solving procedure is similar to that for solving Hartree-Fock problems [10]. In short, the solution algorithm is:

- Obtain a set of guess wave-function u_0 in the multiwavelet basis.
- Iterate until convergence:
 - Compute and diagonalize the Hamiltonian matrix \mathcal{H} to obtain the latest orthogonal wave functions;
 - Compute densities, properties and gradients;
 - Compute the approximation of the Yukawa potential for each eigenvalue λ ;
 - Solve the Lippman-Schwinger equation by convolving with the Green’s function based on Yukawa potential, $u_{n+1} = G \star V u_n$.
- Compute observables.

We carried out benchmark calculations with HFB-AX and MADNESS-HFB for 100 ultra-cold fermionic atoms in elongated traps. The cold Fermions at the unitary limit can be described by the superfluid density functional SLDA [8]. The single-particle Hamiltonian of SLDA can be written as:

$$h(\mathbf{r}) = -\frac{\alpha \hbar^2 \nabla^2}{2m} + \frac{\beta (3\pi^2 \rho(\mathbf{r}))^{2/3}}{2} - \frac{|\Delta(\mathbf{r})|^2}{3\gamma \rho^{2/3}(\mathbf{r})} + V_{ext}(\mathbf{r}), \quad (4)$$

where the deformed external trap potential is:

$$V_{ext}(\mathbf{r}) = V_0 \left[1 - \exp \left(-\frac{\omega^2 (x^2 + y^2 + z^2/\eta^2)}{2V_0} \right) \right]. \quad (5)$$

The densities of spin-up (ρ_\uparrow) and spin-down (ρ_\downarrow) atoms, the pairing densities κ , and pairing gaps Δ can be expressed in terms of the HFB two-component eigenvectors (1):

$$\begin{aligned} \rho_\uparrow(\mathbf{r}) &= \sum_i f_i |u_i(\mathbf{r})|^2, & \rho_\downarrow(\mathbf{r}) &= \sum_i (1 - f_i) |v_i(\mathbf{r})|^2, \\ \kappa(\mathbf{r}) &= \sum_i f_i u_i(\mathbf{r}) v_i^*(\mathbf{r}), & \Delta(\mathbf{r}) &= -g_{eff}(\mathbf{r}) \kappa(\mathbf{r}), \end{aligned} \quad (6)$$

where $\rho = \rho_\uparrow + \rho_\downarrow$ and $g_{eff}(\mathbf{r})$ is the regularized pairing strength. In Eq. (6), $f_i = [1 + \exp(E_i/kT)]^{-1}$ and the temperature is $kT=0.01$. The parameters α , β and γ in Eq.(4) are taken as 1.14, -0.553 and $-1/0.0906$, respectively [8]. We work in trap units for which $\hbar = m = \omega = 1$. The trap aspect ratio η in (5) denotes the trap elongation. In experiments, the adopted optical trap is highly elongated with η up to 50 [17]. The highly elongated trap

is interesting because it provides a connection to quasi one-dimensional systems. In MADNESS-HFB, we employed very large boxes: $(x, y, z) [-100, 100]$ for $\eta = 5$ and $[-160, 160]$ for $\eta = 16$. In HFB-AX, the 2D box sizes are 10.5×35 ($\eta = 5$) and 9.1×70 ($\eta = 16$). The potential depth V_0 is 12 and 10 for $\eta = 5$ and $\eta = 16$, respectively.




Figure 3. Pairing densities $\kappa(x = 0, y = 0, z)$ calculated by MADNESS-HFB and HFB-AX for 100 polarized atoms in an elongated trap with $\eta = 5$ (left) and $\eta = 16$ (right).

Figure 3 displays the calculated pairing densities $\kappa(x = 0, y = 0, z)$ for ultra-cold Fermi gas with polarization of 0.2. The energy cutoff is taken as 6.75 and 5.05 for $\eta = 5$ and 16, respectively. It is seen that the two solvers agree very well at $\eta=5$. At $\eta = 16$, there appears a very small difference between HFB solutions; this may be an indication that an even larger box is needed in HFB-AX. It is to be noted that at $\eta = 16$ there appear coexisting solutions which are close in energy [18, 4]. Figure 4 shows the 3D pairing density distribution computed by MADNESS-HFB for very elongated system with $\eta = 16$. The characteristic transversal oscillations of the pairing field are indicative of the Larkin-Ovchinnikov phase [4].




Figure 4. Pairing density $\kappa(x, y, z)$, calculated by MADNESS-HFB for the elongated trap with $\eta = 16$. The box scale in view is $x [-12, 12]$, $y [-12, 12]$ and $z [-32, 32]$.

In general, the computational effort of 3D HFB calculations is significantly greater than in the 2D case. Thus, the adaptive representation using multiwavelet techniques and parallel

computing are crucial for the success of MADNESS-HFB. To put things in perspective, MADNESS-HFB takes about 5 minutes per iteration on 5000 cores for $\eta = 5$ and $\eta = 16$, while HFB-AX takes about 10 minutes per iteration for $\eta = 5$ on 400 cores and 25 minutes for $\eta = 16$ on 800 cores. From this trend, we conclude that the parallel performance of MADNESS-HFB is superior to HFB-AX in large-box calculations.

4. Conclusions

The coordinate-space HFB solvers HFB-AX and MADNESS-HFB have been developed for large-box calculations of superfluid Fermi systems. The 2D solver HFB-AX has been adapted to the multi-core supercomputers by using a hybrid MPI and OpenMP programming model. The 3D solver MADNESS-HFB is based on multi-resolution multiwavelet techniques and more sophisticated hybrid MPI and pthreads based task parallelism programming methodologies. A scaling benchmark test for 100 ultracold Fermions in elongated traps has been carried out. Developments are underway to efficiently use the next generation of multi-core+GPU architectures for more demanding nuclear problems.

Acknowledgments

This work was supported in part by the U.S. Department of Energy under Contract Nos. DE-FG02-96ER40963 (University of Tennessee) and DE-FC02-09ER41583 (UNEDF SciDAC Collaboration). It is also partially sponsored by the Office of Advanced Scientific Computing Research; U.S. Department of Energy. The work was partially performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725.

Computational resources were provided through an INCITE award “Computational Nuclear Structure” by the National Center for Computational Sciences (NCCS) and National Institute for Computational Sciences (NICS) at Oak Ridge National Laboratory, and the National Energy Research Scientific Computing Center (NERSC).

References

- [1] Pei J C, Nazarewicz W, and Stoitsov M 2009 *Eur. Phys. J. A* **42** 595.
- [2] Pei J C, Stoitsov M V, Fann G I, Nazarewicz W, Schunck N and Xu F R 2008 *Phys. Rev. C* **78** 064306.
- [3] Newton W G and Stone J R 2009 *Phys. Rev. C* **79** 055801.
- [4] Pei J C, Dukelsky J, Nazarewicz W 2010 *Phys. Rev. A* **82** 021603(R).
- [5] <http://icl.cs.utk.edu/magma/>
- [6] Pei J C, Kruppa A T, and Nazarewicz W 2011 *Phys. Rev. C* **84** 024311.
- [7] Stoitsov M, Kortelainen M, Nakatsukasa T, Losa C and Nazarewicz W 2011 *Phys. Rev. C* **84** 041305(R).
- [8] Bulgac A 2007 *Phys. Rev. A* **76** 040502(R).
- [9] Bulgac A and Forbes M M 2008 *arXiv:0808.1436*.
- [10] Fann G I, Pei J, Harrison R J, Jia J, Hill J, Ou M, Nazarewicz W, Shelton W A and Schunck N 2009 *J. Phys. Conf. Ser.* **180** 012080.
- [11] Fann G I, Harrison R J, Beylkin G, Jia J, Hartman-Baker R, Shelton W A and Sugiki S 2007 *J. Phys. Conf. Ser.* **78** 012018.
- [12] <http://code.google.com/p/m-a-d-n-e-s-s/>
- [13] Alpert B 1993 *SIAM J. of Math. Anal.* **24**(1) 246.
- [14] Alpert B, Beylkin G, Gines D and Vozovoi L 2002 *J. of Comp. Phys.* **182** 149.
- [15] Beylkin G and Mohlenkamp M 2002 *Proc. Natl. Acad. Sci. USA.* **99** 10246.
- [16] Harrison R, Fann G, Yanai T, Gan Z and Beylkin G 2004 *J. Chem. Phys.* **23** 11587.
- [17] Partridge G B et al 2006 *Science* **311** 503.
- [18] Baksmaty L O, Lu H, Bolech C J and Pu H 2011 *Phys. Rev. A* **83** 023604.